

CLAIMS

What is claimed is:

- 1 1. A group communication protocol system comprising:
2 a plurality of nodes on a first local area network (LAN), the plurality of nodes
3 logically divided into at least a first group and a second group;
4 a first token to circulate among members of the first group to cause
5 communications among the members of the first group to be serialized;
6 and
7 a second token to circulate among members of the second group to cause
8 communications among the members of the second group to be serialized
9 independent of the first group.
- 1 2. The system of claim 1, wherein at least one member of the first group is also a
2 member of the second group.
- 1 3. The system of claim 1, wherein ownership of the first token is needed before a
2 node can send a message to the first group.
- 1 4. The system of claim 1, wherein the communication among the members of the
2 first group comprises multicast messages.
- 1 5. The system of claim 1, wherein the communication among the members of the
2 second group comprises broadcast messages.
- 1 6. The system of claim 1, wherein the first and second tokens include a sequencing
2 mechanism.

1 7. The system of claim 1 further comprising one or more nodes on a second LAN,
2 wherein the one or more nodes on the second LAN are members of the first group.

1 8. The system of claim 1, wherein the first and second groups comprise replication
2 groups, each including at least one primary and at least one replica.

1 9. A group communication protocol system comprising:
2 a plurality of nodes logically divided into at least a first group and a second group;
3 a first token to circulate among members of the first group to cause
4 communications among the members of the first group to be serialized;
5 and
6 a second token to circulate among members of the second group to cause
7 communications among the members of the second group to be serialized
8 independent of the first group.

1 10. The method of claim 9, wherein ownership of the first token is needed before a
2 node can send a message to the first group.

1 11. The system of claim 9, wherein the communication among the members of the
2 first group comprises unicast messages.

1 12. A method comprising:
2 independently serializing message communication among members of a first
3 group and members of a second group on a local area network (LAN) by
4 circulating a first token among members of the first group; and
5 circulating a second token among members of the second group.

1 13. The method of claim 12, wherein the members of the first group include a primary
2 entity and at least one replica for the primary entity.

1 14. The method of claim 12, wherein the members of the second group include a
2 primary entity and at least one replica for the primary entity.

1 15. The method of claim 12, wherein the first token includes a sequence number.

1 16. The method of claim 15, further comprising:

- 2 a. receiving the first token at a first member of the first group;
- 3 b. incrementing the sequence number
- 4 c. sending a broadcast message to the first group using the sequence number;
- 5 d. repeating b-c for each message, if any, at the head of one or more message
6 queues of the first member that are destined for the first group or until a
7 specified event has occurred; and
- 8 e. passing the first token to the next member of the first group.

1 17. A method comprising:

2 receiving, at a first member of a first group on a local area network (LAN), a first
3 token associated with the first group from another member of the first
4 group on the LAN;
5 incrementing a sequence number associated with the first token;
6 sending a message to the members of the first group using the sequence number
7 associated with the first token;
8 passing the first token to a next member of the first group on the LAN;

receiving, at a member of a second group on the LAN, a second token associated with the second group from another member of the second group on the LAN;
incrementing a sequence number associated with the second token;
sending a message to the members of the second group using the sequence number associated with the second token; and
passing the second token to a next member of the second group on the LAN.

18. The method of claim 17 further comprising:

replacing an all received up-to (aru) field associated with the first token with a lower aru associated with the first member of the first group.

19. The method of claim 17, wherein the sending a message to the members of the first group comprises sending a multicast message.

20. A replication group system comprising:

a first replication group located on a local area network (LAN), the first replication group including a first primary entity and a first group of one or more replica entities wherein members of the first replication group are members of a first group;
a second replication group located on the LAN, the second replication group including a second primary entity and a second group of one or more replica entities wherein members of the second replication group are members of a second group;

an intersection between the first replication group and the second replication group including at least one replica entity that is a member of both the first group and the second group;
a first token circulating among members of the first group causing communications among the members of the first replication group to be ordered; and
a second token circulating among members of the second group causing communications among the members of the second replication group to be ordered independent of the first replication group.

21. The system of claim 20 further comprising:

a first storage area associated with the intersection, comprising serialized messages for the first replication group; and
a second storage area associated with the intersection, comprising serialized messages for the second replication group.

22. The system of claim 20, wherein at least one replica entity in the intersection operates as a warm or cold replica for the first primary entity and a warm or cold replica for the second primary entity.

23. The system of claim 20, wherein at least one replica entity in the intersection operates as a hot replica for the first primary entity and a warm or cold replica for the second primary entity.

1 24. A method comprising:
2 receiving at a first primary entity of a first replication group on a first local area
3 network (LAN) a first token associated with the first replication group, the
4 first token logically imposing a first token ring upon the first replication
5 group, the first replication group including the first primary entity and a
6 first group of one or more replica entities;
7 receiving at the first primary entity a second token associated with a second
8 replication group on the first LAN, the second token logically imposing a
9 second token ring upon the second replication group, the second
10 replication group including a second primary entity and a second group of
11 one or more replica entities;
12 incrementing a sequence number associated with the first token;
13 incrementing a sequence number associated with the second token;
14 sending a message from the first primary entity to the first replication group using
15 the sequence number associated with the first token; and
16 sending a message from the first primary entity to the second replication group
17 using the sequence number associated with the second token.

1 25. The method of claim 24, wherein the first replication group further comprises a
2 replica entity located on a second LAN.

1 26. The method of claim 24, wherein the first and second tokens comprise Totem
2 tokens.

1 27. A method comprising:
2 a step for receiving at a node on a local area network (LAN) a first token
3 associated with a first group on the LAN;
4 a step for sending a first message to the members of the first group using a
5 sequence number associated with the first token;
6 a step for passing the first token on to a next member of the first group;
7 a step for receiving at a second node on the LAN a second token associated with a
8 second group on the LAN;
9 a step for sending a message to the members of the second group using a sequence
10 number associated with the second token; and
11 a step for passing the second token on to a next member of the second group.

1 28. The method of claim 27 further comprising:
2 a step for incrementing the sequence number associated with the first token; and
3 a step for incrementing the sequence number associated with the second token.

1 29. The method of claim 27, wherein the first and second tokens comprise Totem
2 tokens.

1 30. A machine-readable medium having stored thereon data representing sequences of
2 instructions that when executed cause a machine to:
3 receive a first token associated with a first group on a local area network (LAN);
4 send a message to the members of the first group using a sequence number
5 associated with the first token;
6 receive a second token associated with a second group on the LAN; and

7 send a message to the members of the second group using a sequence number
8 associated with the second token.

1 31. The machine-readable medium of claim 30 further including instructions to:
2 increment the sequence number associated with the first token; and
3 pass the first token to a next member of the first group.

1 32. The machine-readable medium of claim 30 wherein the first and second tokens
2 comprise Totem tokens.

1 33. A group communication system comprising:
2 a plurality of nodes on a local area network (LAN) logically divided into a first
3 group and a second group;
4 a first token means, circulating among members of the first group, for serializing
5 multicast communications among the members of the first group; and
6 a second token means, circulating among members of a second group, for
7 serializing multicast communications among the members of the second
8 group independent of the first group.

1 34. The system of claim 33 wherein the first and second token means include a
2 sequence number.

1 35. The system of claim 34, wherein ownership of the first token means is needed
2 before a node can send a message to the first group.

1 36. A method comprising:
2 imposing a first logical token ring on a local area network (LAN) and serializing
3 communications among a first subset of nodes on the LAN by causing a
4 first token to be circulated among the first subset of nodes; and
5 imposing a second logical token ring on the LAN and serializing communications
6 among a second subset of nodes on the LAN by causing a second token to
7 be circulated among the second subset of nodes.

1 37. The method of claim 36, wherein the first and second tokens comprise Totem
2 tokens.

1 38. The method of claim 36, wherein ownership of the first token is needed before a
2 node can send a message to the first subset of nodes.